

Vantalect



AWS COST · SAMPLE DELIVERABLE

Week-Six Savings Measurement

Week 6 of 6: Measure · prepared for Copperwick Therapeutics, Inc.

Prepared by Charles Sieg · charles.sieg@vantalect.com · vantalect.com

April 1, 2026 · AWS Cost Analysis & Optimization: The Optimization (six weeks)

SAMPLE DELIVERABLE

This is a sample deliverable. Copperwick Therapeutics, Inc. is a fictional company: its people, programs, products, accounts, and figures are invented for this sample, and any resemblance to a real company, person, program, or medicine is coincidental. AWS prices are list prices from the AWS Price List API retrieved on October 13, 2026, and from the public AWS Price List bulk API on October 17, 2026; Spot prices are 30-day means for September 14 to October 14, 2026.

Contents

Result	2
The method	4
The query for one change	7
Measured against estimated	10
Every change	12
Where the measurement fell short, and where it did not	15
The gate	17
Robustness	18
The private pricing agreement	19
One-time charges.....	19
In place, not yet measurable	20
Approved, then declined	20
What is not counted	21
Re-running this measurement	22

Result

The 22 changes in place cut **\$2,401,346 a year** from Copperwick's AWS bill, **13.1% of the \$18,315,196 in-scope annual bill, measured on Copperwick's Cost and Usage Report and net of the 6% private pricing discount**. That is 95.1% of the \$2,526,309 the Analysis estimated for the same changes. Counting the one change approved at the readout and then reversed, as the engagement's terms define it, the figure for the guarantee is **\$2,436,540 a year**, 13.3% of the bill, against a gate of \$105,000. The gate is met with \$2,331,540 to spare, and the second invoice, \$17,500, is issued with this report.

Net means after the discount, because the discount is what Copperwick's bill actually falls by. The same measurement at list prices is \$2,554,624; list figures appear in this report only as context and are always labeled "at list." The share of the bill is the same on either basis.

Without the reversed change the margin over the gate is \$2,296,346, so the gate does not depend on it. The \$35,000 fee is everything Copperwick paid: the Analysis was the engagement's first two weeks, not a second fee. The measured savings repay it in about 5 days,


and the first year nets \$2,366,346 after it. The counted figure decides the gate and nothing else. The one-time charges of the changes, \$177.66 net, are listed in [the one-time charges](#) and never annualized; with them deducted the counted figure is \$2,436,362.

The result against the gate. The full computation is in [the gate computation](#).

The result, net of the private pricing discount	Amount
Measured savings, 22 changes in place, annualized	\$2,401,346
Share of the in-scope bill (\$18,315,196 a year, net)	13.1%
Counted for the gate, with L-4 at its estimate	\$2,436,540
The gate: three times the fee of \$35,000	\$105,000
Margin over the gate	\$2,331,540
Payback on the fee, from the measured savings	About 5 days
First-year net, measured savings less the fee	\$2,366,346
For context, measured savings at list	\$2,554,624

The second invoice

The Optimization's second \$17,500 is invoiced only if the annualized savings in place at the end of week six reach three times the fee, \$105,000. They reach \$2,436,540, so the invoice is issued with this report. If you disagree with any figure here, the queries in [the queries](#) reproduce it from Copperwick's own data.



From what the Analysis identified to what the Cost and Usage Report measured, net of the discount. Every step is a row in [the gate computation](#).

The method

The terms of the engagement say how the second half of the fee is decided:

The Optimization is invoiced in full only if the annualized savings in place at the end of week six are at least three times the fee. Savings are measured change by change on the client's Cost and Usage Report: the daily cost of each changed resource or commitment in the week before the change against the week after it, annualized. A change the client approved at the week-two readout and later declines counts at the savings the roadmap estimated for it.

Each step below applies that sentence to Copperwick's bill, the same way to every change. Every figure is net of the 6% private pricing discount: the discount is a flat percentage on eligible usage, Savings Plan fees included, so a saving of \$1.00 at list is \$0.94 on Copperwick's invoice.

- 1. The cost.** Every change except B-1 and C-1 is measured on **public On-Demand cost** (pricing_public_on_demand_cost on usage a plan covered, the unblended cost elsewhere), and then multiplied by 0.94, one minus the discount. The source is the CUR 2.0 data export in Parquet, queried with Athena; one query per change, the same for every change. Public On-Demand cost is the price before any plan or discount, so the existing three-year Compute Savings Plan (\$150 an hour, in place until July 31, 2027), C-1, and the private pricing discount are each counted exactly once: the plans are not in any other change's cost, and the discount is applied once, as a factor. Savings Plan sharing is restricted to the 48 in-scope accounts (a cost category) (since January 12), so the plans cover usage in those accounts only. The existing plan stays fully used after every change, because the always-on floor of \$395.50 an hour after every deferred change is above the \$363.83 an hour the two plans cover together. The expression is in [the queries](#).
- 2. Spot (B-1)** is measured at the **Spot price paid** (line_item_unblended_cost on the Spot usage lines), also times 0.94. Public On-Demand cost on a Spot line shows the On-Demand equivalent and would hide the saving.
- 3. Volume-driven changes (B-1 and N-1a)** are normalized. Genomics spend moves with lab throughput, so for these two the measurement compares the cost per 30x whole-genome equivalent processed, before and after, times the 90-day average volume. Without that, a slow lab week would read as a saving and a busy one as a loss. The query is in [the per-genome query](#).
- 4. The window.** For each changed resource, take its cost for each of the 7 days before the change and the 7 days after it, and average each week into a daily mean. The day of the change is excluded from both, because a change lands partway through its day. The daily saving is the mean before minus the mean after, and the annual saving is that difference times 365. Each measurement is that difference of two weekly means, from the finalized daily rows, not a percentage of the estimate.
- 5. What is left out.** These line types are excluded from every change's cost: SavingsPlanRecurringFee, SavingsPlanNegation, RIFee, Fee, Tax, Credit, Refund, EdpDiscount, PrivateRateDiscount. Recurring fees, negations, taxes, credits, and refunds are not usage a change moves, and the private pricing discount arrives as its own negative lines rather than on the usage lines, so a per-resource query cannot see it; it is applied as the factor in step 1.
- 6. The lag.** The Cost and Usage Report finalizes a day about 48 hours later. This measurement ran on Wednesday, April 1, so it uses data through March 30. A change needs 7 full days after it, so only changes made by March 22 could be measured. Every counted change was in place by March 19, with the first on March 10, which is why week 5 carries no savings by design.

7. **The Savings Plan (C-1)** is measured by its ARN, on the usage it covered in the 48 in-scope accounts: its fee, the On-Demand cost of that covered usage, and its utilization, not by a resource ID ([the Savings Plan query](#)).
8. **S3 changes (W-5, S3-1, S3-2)**. S3 is billed per bucket, and W-5, S3-1, S3-2, and S3-4 changed the same lake buckets from March 11 to 13, so each S3 change is attributed by prefix and storage class: S3 Inventory and S3 Storage Lens give the bytes in each class under each change's prefixes per day, times the published price, reconciled to the bucket's CUR total. A Batch job tars each run with a checksum manifest, verifies it, and transitions it; after the owner's sign-off it deletes the originals by version ID, so no noncurrent copy stays billed in the versioned bucket. S3-1's 97% has its own reason: checksum manifests and tars still in verification, 10.5 TB, billed at Standard in the window.
9. **One-time charges** are listed beside the result and never annualized into it.
10. **Nothing is counted that is not in place**. Deferred and declined changes count for nothing, and so does a change that is in place but cannot yet be measured ([the section on S3-4](#)). The one exception is the rule the terms state for a change approved at the readout and declined later ([the rule for a reversed change](#)).

How each counted change was measured. B-1 and N-1a are also normalized per 30x genome equivalent, and W-5, S3-1, S3-2 are attributed by prefix and storage class (step 8).

Measured by	ChangesWhich
Public On-Demand cost, times the discount's complement	19D-1, E-1, E-2a, EF-1, F-1, L-1, L-2, L-3, R-1, S-1, S-2, S-3a, S3-1, S3-2, W-1, W-2, W-3, W-4, W-5
Spot price paid, times the discount's complement	1B-1
Public On-Demand cost per 30x genome equivalent	1N-1a
The Savings Plan's ARN	1C-1
Total	22

A worked example

W-1, deleting the unattached gp2 volumes, was made on Tuesday, March 10, after volumes carrying 14 TB were snapshotted. The seven daily rows in the seven days before (March 3 to March 9) average \$284.81 a day for the 90 TB of volumes, net of the discount. The seven daily rows in the seven days after (March 11 to March 17) average \$15.19 a day for the retained snapshots, which bill 9.6 TB of written blocks, not the 14 TB of volumes they came from. The daily saving is \$269.62; times 365 is \$98,412. The Analysis estimated \$95,871 net, pricing all 14 TB, so W-1 measures at 102.7% of its estimate, above it.

The query for one change

Anyone on Copperwick's team can run the measurement for a single change in Athena against the CUR 2.0 export. The query below is W-1's; the volume IDs come from the waste register, where each deletion has a signed-off line. Replace the register filter, the dates, and the billing periods for another change. The window is summed and divided by 7, not averaged: a deleted resource has no rows after the change, and a day with no rows is a day with no cost, so an average would return nothing for the changes that save the most.

SQL

```
WITH daily AS (
  SELECT date(line_item_usage_start_date) AS day,
         sum(CASE line_item_line_item_type
              WHEN 'SavingsPlanCoveredUsage' THEN pricing_public_on_demand_cost
              WHEN 'DiscountedUsage' THEN pricing_public_on_demand_cost
              ELSE line_item_unblended_cost
              END) * 0.94 AS cost
  FROM cur2.copperwick_cur
  WHERE line_item_resource_id IN (
    SELECT resource_id FROM cwdata.waste_register WHERE finding = 'W-1'
  )
  AND line_item_line_item_type NOT IN (
    'SavingsPlanRecurringFee',
    'SavingsPlanNegation',
    'RIFee',
    'Fee',
    'Tax',
    'Credit',
    'Refund',
    'EdpDiscount',
    'PrivateRateDiscount'
  )
  AND billing_period IN ('2026-03')
  GROUP BY 1
), windows AS (
  SELECT sum(CASE WHEN day BETWEEN date '2026-03-03' AND date '2026-03-09' THEN cost ELSE 0
  END) / 7.0 AS daily_before,
```

```

        sum(CASE WHEN day BETWEEN date '2026-03-11' AND date '2026-03-17' THEN cost ELSE 0
END) / 7.0 AS daily_after
    FROM daily
)
SELECT daily_before,
       daily_after,
       (daily_before - daily_after) * 365 AS annual_saving
FROM windows;

```

For a change that replaces resources, filter on tags in place of resource IDs; user-defined tag keys appear in the `resource_tags` map with a `user_prefix`, as in `element_at(resource_tags, 'user_workload') = 'genomics-batch'`. Because W-1 snapshots 14 TB of volumes first, its after-window is not zero: the snapshots, which bill 9.6 TB of written blocks, are in the retained resource list.

B-1 and N-1a: cost per genome

Genomics Batch and the S3 traffic of the research VPCs move with the number of genomes the lab processes. B-1 (Spot, priced at the Spot price paid) and N-1a (gateway endpoints that take S3 reads off the NAT gateways) are therefore measured as a unit cost: the cost per 30x whole-genome equivalent before and after, times the 90-day average daily volume, times 365. The denominator is the pipeline's own completion log, loaded into Athena as `cwdata.genomics_completed` (one row per day: `day` and `genomes_30x`).

SQL

```
WITH daily AS (
  SELECT date(line_item_usage_start_date) AS day,
         sum(CASE
              WHEN line_item_usage_type LIKE '%SpotUsage%' THEN line_item_unblended_cost
              WHEN line_item_line_item_type IN ('SavingsPlanCoveredUsage',
'DiscountedUsage') THEN pricing_public_on_demand_cost
              ELSE line_item_unblended_cost
            END) * 0.94 AS cost
  FROM cur2.copperwick_cur
  WHERE element_at(resource_tags, 'user_Workload') = 'genomics-batch'
     AND line_item_line_item_type NOT IN ('SavingsPlanRecurringFee', 'SavingsPlanNegation',
'RIFee', 'Fee', 'Tax', 'Credit', 'Refund', 'EdpDiscount', 'PrivateRateDiscount')
     AND billing_period IN ('<yyyy-mm-before>', '<yyyy-mm-after>')
  GROUP BY 1
), unit AS (
  SELECT d.day, d.cost, g.genomes_30x
  FROM daily d JOIN cwdata.genomics_completed g ON g.day = d.day
), windows AS (
  SELECT sum(CASE WHEN day BETWEEN date '<before-start>' AND date '<before-end>' THEN cost END)
        / sum(CASE WHEN day BETWEEN date '<before-start>' AND date '<before-end>' THEN
genomes_30x END) AS cost_per_genome_before,
        sum(CASE WHEN day BETWEEN date '<after-start>' AND date '<after-end>' THEN cost END)
        / sum(CASE WHEN day BETWEEN date '<after-start>' AND date '<after-end>' THEN
genomes_30x END) AS cost_per_genome_after,
        avg(genomes_30x) FILTER (WHERE day > date_add('day', -90, date '<change-date>') AND
day < date '<change-date>') AS genomes_per_day_90
  FROM unit
)
SELECT cost_per_genome_before,
       cost_per_genome_after,
       (cost_per_genome_before - cost_per_genome_after) * genomes_per_day_90 * 365 AS
annual_saving
FROM windows;
```

For N-1a, use the same shape on the NAT gateway lines of the research VPCs in place of the Batch tag. The third weekly sequencing run kept growing the genomics line during the engagement, and this normalization is what keeps that growth out of both changes.

The Savings Plan, by its ARN

C-1 is measured on the plan, not on a resource. The first query reads the On-Demand cost of the usage the plan covered and the plan's fee; the utilization comes from the recurring-fee lines that the cost measurement leaves out. The plan costs \$50.00 an hour, \$1,200.00 a day at list. Savings Plan sharing is restricted to the 48 in-scope accounts (a cost category). C-1 was bought in cw-cicd, a member account inside that group (a payer account cannot belong to one), as the existing plan was in 2024, so it covers usage in the 48 in-scope accounts only, and the query filters covered usage to them. In the measurement window it covered \$1,652.64 a day of On-Demand usage in those accounts, so its \$1,200.00 fee bought an effective discount of 27.4%, a saving of \$452.64 a day at list, \$165,213 a year, or \$155,300 net of the discount, at 100% utilization. The

Analysis priced it at \$1,633.35 a day of coverage, the one-year m6i rate, 26.5% off. It was bought on March 19, so there is no plan before that day: the saving is the covered usage less the fee over the seven days after it, March 20 to March 26.

SQL

```
SELECT date(line_item_usage_start_date) AS day,
       sum(CASE WHEN line_item_line_item_type = 'SavingsPlanCoveredUsage' THEN
pricing_public_on_demand_cost ELSE 0 END) AS covered_on_demand,
       sum(CASE WHEN line_item_line_item_type = 'SavingsPlanRecurringFee' THEN
savings_plan_used_commitment ELSE 0 END) AS fee_for_used_commitment,
       sum(CASE WHEN line_item_line_item_type = 'SavingsPlanRecurringFee' THEN
savings_plan_used_commitment ELSE 0 END)
       / nullif(sum(CASE WHEN line_item_line_item_type = 'SavingsPlanRecurringFee' THEN
savings_plan_total_commitment_to_date ELSE 0 END), 0) AS utilization
FROM cur2.copperwick_cur
WHERE savings_plan_savings_plan_a_r_n = '<savings-plan-arn>'
      AND line_item_usage_account_id IN (SELECT account_id FROM cwdata.in_scope_accounts) -- the
sharing group: the 48 in-scope accounts
      AND billing_period IN ('2026-03')
GROUP BY 1
ORDER BY 1;
-- daily saving, at list = covered_on_demand - fee_for_used_commitment; times 365, then times
0.94 for the net figure
```


Two limits are worth stating. The measurement is of cost, not usage, so a workload that grew or shrank between the two weeks moves a single change's figure for reasons that are not the change; across 22 changes those effects average out, and for any one change they do not, which is why the next sections show the estimate beside the measurement. And the report is provisional until AWS finalizes the month's invoice: we re-run the measurement at the first monthly review and say so if it moves.

Measured against estimated

Overall the measurement is 95.1% of the estimates. We expected the measurements to fall a little under the estimates: an estimate is made from the 30 days before the Analysis, and a measurement is made from the week after a change, when some usage has moved. Of the 22 changes, 6 measured at their estimates to the dollar, 4 measured below 90%, 10 measured between 90.8% and 99.8%, and 2 measured above their estimates (C-1 and W-1). \$134,122 of net savings fell short of the estimates on 14 changes, and the \$9,159 of extra saving on C-1 and W-1 offsets part of it, which is the \$124,963 difference between the \$2,526,309 estimated and the \$2,401,346 measured.

Every counted change, by category, estimated and measured, annualized and net of the discount.

Category	Changes	Estimated	Measured	Measured vs estimated
Waste (W-1, W-2, W-3, W-4, W-5)	5	\$644,816	\$618,335	95.9%
Schedules (S-1, S-2, S-3a)	3	\$335,838	\$289,870	86.3%
File systems (F-1)	1	\$255,040	\$255,040	100.0%
Batch and Spot (B-1)	1	\$238,111	\$216,229	90.8%
Data lake (S3-1, S3-2)	2	\$218,155	\$210,502	96.5%
Storage (E-1, E-2a, EF-1)	3	\$198,673	\$185,613	93.4%
Logging and governance (L-1, L-2, L-3)	3	\$188,613	\$184,306	97.7%
Commitments (C-1)	1	\$148,682	\$155,300	104.5%
Network (N-1a)	1	\$135,143	\$122,913	91.0%
Rightsizing (R-1)	1	\$113,832	\$113,832	100.0%
Databases (D-1)	1	\$49,406	\$49,406	100.0%
Total	22	\$2,526,309	\$2,401,346	95.1%



Measured savings by category. The ordinary research-estate categories carry most of the result.

Every change

The 22 changes in place, in the order they were made, each with the seven full days of finalized data after it and the seven before it, the change day excluded. The list columns are context; the net column is the measurement. Basis: public On-Demand cost, except B-1 (Spot price paid), N-1a (per genome), and C-1 (its ARN).

ID	Change	Applied	Estimate at list	Measured at list	Measured, net	vsWhy it measured as it did est.
D-1	Multi-AZ off on development databases	Mar 10	\$52,560	\$52,560	\$49,406	100%Six instances before and after; only the standby went away, and storage bills the same
E-1	gp2 to gp3, research and platform	Mar 10	\$103,219	\$98,875	\$92,943	96%38 volumes above 1 TiB needed 58,000 IOPS and 1,800 MiB/s above gp3's included baseline, at \$0.005 an IOPS-month and \$0.04 a MiB/s-month; the estimate priced storage only
L-1	Batch job logs to Infrequent Access class	Mar 10	\$104,448	\$99,947	\$93,950	96%Two pipelines still wrote to the old groups for two days; and day-to-day usage in the two windows took 0.9 points from the 96.6% those alone give
L-2	Research log retention to 90 days	Mar 10	\$61,932	\$61,932	\$58,216	100%The expired bytes left in the first days, and what stays (90 days of the other research groups, 30 days of Batch job logs) is level from day to day
N-1a	S3 gateway endpoints, research VPCs	Mar 10	\$143,770	\$130,759	\$122,913	91%Two research VPCs whose route tables sit outside the platform template kept 21 TB a month of S3 traffic on NAT until their routes were added; and day-to-day usage in the two windows took 1.0 points from the 91.9% those alone give
S-2	Development accounts on a weekday schedule	Mar 10	\$139,068	\$130,593	\$122,758	94%Two build hosts carry exemption tags and ran around the clock
W-1	Unattached gp2 volumes deleted	Mar 10	\$101,990	\$104,694	\$98,412	103%The 14 TB of volumes owners kept snapshotted to 9.6 TB of written blocks; a snapshot bills its blocks, not its volume's size, and the estimate priced all 14 TB
W-2	Notebooks stop after 60 idle minutes	Mar 10	\$193,778	\$178,790	\$168,063	92%Two notebooks carry an exemption tag for long-running training and billed 34% of hours, not 22%
EF-1	EFS lifecycle to Infrequent Access	Mar 11	\$94,618	\$85,068	\$79,964	90%64% of the data moved to Infrequent Access, not 70%, and the measurement includes the read charge on 6.5 TB a month; and day-to-day usage in the two windows took 0.7 points from the 90.6% those alone give

ID	Change	Applied	Estimate at list	Measured at list	Measured, net	vsWhy it measured as it did est.
L-3	Config daily recording, high-churn accounts	Mar 11	\$34,272	\$34,192	\$32,140	100%Instance churn varies from week to week
S-1	Workstation idle-stop and schedule	Mar 11	\$125,078	\$101,545	\$95,452	81%Six scientists overrode auto-stop during a structure campaign, about 272 more hours each a month
W-4	Orphaned snapshots deleted	Mar 11	\$92,160	\$78,397	\$73,694	85%22.4 TB of the deleted snapshots' blocks were shared with retained snapshots and still bill on them
W-5	Noncurrent versions and multipart uploads expire	Mar 11	\$105,696	\$103,572	\$97,358	98%8.3 TB of multipart uploads were still in flight and not yet 30 days old
E-2a	gp2 to gp3, validated, no freeze	Mar 12	\$13,517	\$13,517	\$12,706	100%No volume was above 1 TiB, so none needed IOPS or throughput above gp3's included baseline
R-1	Data platform services to m6i.2xlarge	Mar 12	\$121,098	\$121,098	\$113,832	100%Thirty-six instances before and after, around the clock; only the instance type changed
S3-2	Cryo-EM movies to Glacier Flexible Retrieval	Mar 12	\$131,366	\$125,807	\$118,258	96%26 TB of movies carry object dates newer than their acquisition dates, so they have not reached 180 days
W-3	One idle p4d.24xlarge released	Mar 12	\$192,349	\$192,349	\$180,808	100%One instance released, and nothing billed in its place
S3-1	Raw run folders tarred to Deep Archive	Mar 13	\$100,714	\$98,132	\$92,244	97%Checksum manifests and tars still in verification, 10.5 TB, billed at Standard in the window
B-1	Genomics Batch: 40% to Spot	Mar 16	\$253,310	\$230,031	\$216,229	91%Spot prices in the window ran about 6% above the 30-day mean the estimate used, and retries were 15% of Spot hours, not 12%; and day-to-day usage in the two windows took 1.1 points from the 91.9% those alone give
F-1	Genomics Lustre resized and re-tiered	Mar 17	\$271,319	\$271,319	\$255,040	100%A Lustre file system bills its provisioned size and throughput, which did not change from day to day

ID	Change	Applied	Estimate at list	Measured at list	Measured, net	vsWhy it measured as it did est.
C-1	Compute Savings Plan, one year	Mar 19	\$158,172	\$165,213	\$155,300	104%Priced at m6i's one-year rate (26.5% off), the lowest it covers; it applied first to c6i usage at 27.7% off and the rest of each hour to m6i, so its \$1,200 a day covered \$1,652.64 of On-Demand usage, 27.4% off
S-3a	GMP val environments on a schedule	Mar 19	\$93,128	\$76,234	\$71,660	82%An executed LIMS operational qualification protocol held its val environment on for three nights of the seven
Total			\$2,687,562	\$2,554,624	\$2,401,346	95.1%

Where the measurement fell short, and where it did not

Every change that measured away from its estimate has a reason, listed in the last column above and discussed here, because a measurement that reports only its wins is not a measurement.

Four changes measured below 90%.

The changes that measured under 90% of their estimates, net of the discount.

ID	Change	Estimated, net	Measured, net	Short by	vs estimate
S-1	Workstation idle-stop and schedule	\$117,574	\$95,452	\$22,122	81.2%
S-3a	GMP val environments on a schedule	\$87,540	\$71,660	\$15,880	81.9%
W-4	Orphaned snapshots deleted	\$86,630	\$73,694	\$12,936	85.1%
EF-1	EFS lifecycle to Infrequent Access	\$88,941	\$79,964	\$8,977	89.9%

S-1, Workstation idle-stop and schedule (81.2%). Six scientists overrode auto-stop during a structure campaign, about 272 more hours each a month. The schedule works; the override tag is the control doing its job, and a structure campaign is the reason to have one. We did not tighten it. The monthly review reads the override count, and the saving is \$22,122 a year less than estimated until campaigns end.

S-3a, GMP val environments on a schedule (81.9%). An executed LIMS operational qualification protocol held its val environment on for three nights of the seven. The saving is real and \$15,880 a year smaller than estimated.

W-4, Orphaned snapshots deleted (85.1%). 22.4 TB of the deleted snapshots' blocks were shared with retained snapshots and still bill on them. Deleting a snapshot frees only the blocks no other snapshot references, so the saving is smaller than the sum of the deleted snapshots' billed size.

EF-1, EFS lifecycle to Infrequent Access (89.9%). 64% of the data moved to Infrequent Access, not 70%, and the measurement includes the read charge on 6.5 TB a month; and day-to-day usage in the two windows took 0.7 points from the 90.6% those alone give. The estimate counted the move to Infrequent Access alone; the measurement counts the read charge on files that are read afterward.

The ten changes between 90.8% and 99.8%. Each has its own reason, from its own daily rows; none is noise we could not name, and none changes the result. B-1 and N-1a are also normalized per genome ([the per-genome query](#)), so the lab's volume is out of both.

The changes that measured between 90.8% and 99.8% of their estimates, each with its reason.

ID	Change	vsWhy estimate
B-1	Genomics Batch: 40% to Spot	90.8%Spot prices in the window ran about 6% above the 30-day mean the estimate used, and retries were 15% of Spot hours, not 12%; and day-to-day usage in the two windows took 1.1 points from the 91.9% those alone give
N-1a	S3 gateway endpoints, research VPCs	91.0%Two research VPCs whose route tables sit outside the platform template kept 21 TB a month of S3 traffic on NAT until their routes were added; and day-to-day usage in the two windows took 1.0 points from the 91.9% those alone give
W-2	Notebooks stop after 60 idle minutes	92.3%Two notebooks carry an exemption tag for long-running training and billed 34% of hours, not 22%
S-2	Development accounts on a weekday schedule	93.9%Two build hosts carry exemption tags and ran around the clock
L-1	Batch job logs to Infrequent Access class	95.7%Two pipelines still wrote to the old groups for two days; and day-to-day usage in the two windows took 0.9 points from the 96.6% those alone give
S3-2	Cryo-EM movies to Glacier Flexible Retrieval	95.8%26 TB of movies carry object dates newer than their acquisition dates, so they have not reached 180 days
E-1	gp2 to gp3, research and platform	95.8%38 volumes above 1 TiB needed 58,000 IOPS and 1,800 MiB/s above gp3's included baseline, at \$0.005 an IOPS-month and \$0.04 a MiB/s-month; the estimate priced storage only
S3-1	Raw run folders tarred to Deep Archive	97.4%Checksum manifests and tars still in verification, 10.5 TB, billed at Standard in the window
W-5	Noncurrent versions and multipart uploads expire	98.0%8.3 TB of multipart uploads were still in flight and not yet 30 days old
L-3	Config daily recording, high-churn accounts	99.8%Instance churn varies from week to week

Two changes measured above their estimates.

W-1, Unattached gp2 volumes deleted (102.7%). The 14 TB of volumes owners kept snapshotted to 9.6 TB of written blocks; a snapshot bills its blocks, not its volume's size, and the estimate priced all 14 TB. It beat its estimate by \$2,541 a year.

C-1, Compute Savings Plan, one year (104.5%). Priced at m6i's one-year rate (26.5% off), the lowest it covers; it applied first to c6i usage at 27.7% off and the rest of each hour to m6i, so its \$1,200 a day covered \$1,652.64 of On-Demand usage, 27.4% off. It beat its estimate by \$6,618 a year.

Together they are the only changes that beat their estimates, by \$9,159 a year.

The six changes that measured at their estimates to the dollar (D-1, E-2a, F-1, L-2, R-1, W-3). Each matched for its own reason: D-1, six instances before and after; only the standby went away, and storage bills the same; L-2, the expired bytes left in the first days, and what stays (90 days of the other research groups, 30 days of Batch job logs) is level from day to day; E-2a, no volume was above 1 TiB, so none needed IOPS or throughput above gp3's included baseline; R-1, thirty-six instances before and after, around the clock; only the instance type changed; W-3, one instance released, and nothing billed in its place; F-1, a Lustre file system bills its provisioned size and throughput, which did not change from day to day. These rows show no scatter, and they prove less than the changes whose quantity moves with traffic: the measurement confirms the price and the date, not a forecast of usage.

The gate

The gate computation. The second invoice is issued because the counted figure is above the gate, with or without the one-time charges. Identified is at the estimates, so it includes the \$124,963 by which the changes in place measured below theirs; [the reconciliation table](#) adds it up.

The guarantee's computation, net of the private pricing discount	Amount
Measured savings, 22 changes in place, annualized	\$2,401,346
Approved at the readout, then reversed, counted at the estimate (L-4)	\$35,194
Counted for the gate	\$2,436,540
The gate: three times the fee of \$35,000	\$105,000
Margin over the gate	\$2,331,540
One-time charges of the counted changes, listed in the one-time charges	\$177.66
Counted, net of the one-time charges	\$2,436,362
Not counted: in place, savings begin as objects age (S3-4)	\$119,686

The guarantee's computation, net of the private pricing discount	Amount
Not counted: approved, deferred (13 changes)	\$1,192,737
Not counted: declined at the readout (C-2, D-2, R-5, B-2)	\$436,118
Identified by the Analysis	\$4,310,044

Stated once, plainly: the counted \$2,436,540 is 69.6 times the \$35,000 fee, which is 23.2 times the \$105,000 gate. A multiple of a fixed fee at a bill this size is not the number to judge. The numbers to judge are the dollar margin of \$2,331,540 and the share of the bill, 13.1% measured against 23.5% identified.

The identified savings reconcile to the dollar. Every dollar the Analysis identified is in exactly one of five places:

In place plus deferred plus declined plus not yet measurable plus approved-then-reversed equals identified; the gap is \$0. At the estimates, net of the discount.

Where the identified savings are	Changes	Annualized
In place at the Analysis's estimates	22: B-1, C-1, D-1, E-1, E-2a, EF-1, F-1, L-1, L-2, L-3, N-1a, R-1, S-1, S-2, S-3a, S3-1, S3-2, W-1, W-2, W-3, W-4, W-5	\$2,526,309
Approved, deferred past the engagement	13: B-1b, B-3, E-2b, F-3, G-1, N-1b, R-2, R-3, R-4, S-3b, S3-1b, S3-3, S3-5	\$1,192,737
Declined at the readout	4: B-2, C-2, D-2, R-5	\$436,118
In place, not yet measurable	1: S3-4	\$119,686
Approved at the readout, then declined	1: L-4	\$35,194
Identified by the Analysis		\$4,310,044

Robustness

The gate does not rest on the soft numbers. The 17 high-confidence changes in place (C-1, D-1, E-1, E-2a, F-1, L-1, L-2, L-3, N-1a, R-1, S-2, S-3a, S3-1, W-1, W-2, W-3, W-5) total \$1,871,568 a year at their estimates, net. Measured at only 90% of their estimates they would be \$1,684,411, which clears the gate by \$1,579,411 on those changes alone, without the medium-confidence changes (B-1, EF-1, S-1, S3-2, W-4) and without L-4.

The cushion is wide in dollars. The counted figure is \$2,331,540 above the gate, so the measured savings would have to be \$69,806 instead of \$2,401,346, a miss of 97.1%, before the gate failed. The week-two stop rule exists for the case where the roadmap is too small: the roadmap identified \$4,310,044, and the engagement continued.

The private pricing agreement

Copperwick has a private pricing agreement with AWS: a three-year term from July 1, 2024 to June 30, 2027, a commitment of \$27,000,000 a year in spend net of the discount, and a flat 6% discount on eligible usage, Savings Plan fees included. Every figure in this report is net of that discount, and the CFO’s question is whether cutting spend risks a shortfall against the commitment. It does not.

Annual spend of all 212 accounts, at list less the discount, against the commitment. The \$1,476,000 a month outside the 48 accounts in scope is unchanged throughout.

The whole organization, a year	Spend, net	Above the \$27,000,000 commitment
Today: about \$3,099,688 a month at list, \$37,196,251 a year at list	\$34,964,476	\$7,964,476
After the 22 measured changes	\$32,563,130	\$5,563,130
After the whole identified roadmap, declined changes included	\$30,654,432	\$3,654,432

All three rows stay above the commitment, and the third row assumes every identified change lands, which the declined ones will not, so it is the most conservative. The margin does narrow: the commitment is 77.2% of today’s net spend and 88.1% of the spend after the roadmap, before planned growth. Planned 2026 growth in the in-scope accounts is about \$260,000 a month by March 2027, \$2,932,800 a year net, which would put the post-roadmap figure at \$33,587,232, \$6,587,232 above the commitment. The savings inform the 2027 renewal: Copperwick can commit to a lower annual figure, or the same figure at a deeper discount. That value is real and is not counted. The handover report sets out the renewal plan.

One-time charges

Some changes cost money once. The charges below are not netted into the savings above, and they are never annualized; the gate was checked both ways.

One-time charges of the changes in place. Each is repaid in the days of that change’s measured saving shown.

ID	Change	At list	Net of theWhat it is discount	Repaid in
S3-1	Raw run folders tarred to Deep Archive	\$141.00	\$132.54The tar job's compute and requests	Within a day
S3-2	Cryo-EM movies to Glacier Flexible Retrieval	\$8.00	\$7.52Lifecycle transition requests	Within a day
F-1	Genomics Lustre resized and re-tiered	\$40.00	\$37.60The repository import requests	Within a day
Total		\$189.00	\$177.66	

The total is 0.01% of one year’s measured savings. S3-4’s own transition charge, \$340 at list, is not in the table because S3-4 is not counted.

In place, not yet measurable

S3-4, Analysis results to Intelligent-Tiering, is in place and counts for \$0 toward the gate. It was applied on March 11. Objects move tiers only after 30 days without access, so nothing can be measured before week six; the first measurement is in May. Its estimate is \$119,686 a year, net, and it is reported separately: it is in neither the measured figure nor the counted one. Its first measurement window runs from Monday, May 11, to May 17; the result is ready May 19, and the monthly cost review on June 11 reads it. That window measures the Infrequent Access tier only: the first objects reach Archive Instant Access on June 9, so the estimate is not measured in full until a later window.

Approved, then declined

The rule

A change approved at the week-two readout and later declined counts at the savings the roadmap estimated for it. The rule exists so that the guarantee cannot be defeated by a change being approved and then not made, and so that Copperwick is not penalized for an owner’s reasonable second thought.

It applies to one change this engagement.

The change approved at the readout and reversed afterward, counted at its estimate. It is the only reason the counted figure exceeds the measured one.

ID	Change	Decided by	Estimated annual savings, net	Reason
L-4	CloudTrail read events dropped on research buckets	Rosalind Tran	\$35,194	Rosalind Tran reversed it on March 11, before it was applied: the read events on research buckets feed the insider-risk detection for sequence designs.

We did not push. It is reported separately from the measured figure, at the estimate, because there is no measurement of a change that was never made.

What is not counted

What the result deliberately leaves out. Savings are net of the private pricing discount; the Support fall is the same dollars at the Support rate on top.

Not counted	Amount, aWhy year
S3-4, in place but not yet measurable	\$119,686Objects move tiers after 30 days without access; first measurement window opens May 11
Approved and deferred (13 changes)	\$1,192,737Waiting on freezes, change control, revalidation, a disaster-recovery test, or a campaign; each has an owner and a date in the handover report
Declined at the readout (4 changes)	\$436,118The owners' reasons are in the handover report with what would reopen each
The fall in Enterprise Support	AboutSupport is a percentage of the whole organization's usage (3% \$72,040on the top tier), so any saving lowers it. It bills monthly, so it is in no daily series and no per-change measurement
The Oracle processor licenses R-3 frees	Not pricedR-3 is deferred; the licenses are Copperwick's own agreement with Oracle (the LIMS database is bring-your-own-license), not an AWS charge
Growth during the engagement	Not a saving orThe third weekly sequencing run kept growing the genomics line; a lossthe per-genome measurement absorbs it for B-1 and N-1a
Cost avoidance on planned growth	\$250,000 toNot measured, not counted toward the gate, and not a saving on \$410,000the bill

The avoidance range comes from the week-five guardrails applied to Copperwick's planned growth of about \$260,000 a month: \$20,800 to \$33,800 a month at the March 2027 run rate. It is **not measured, not counted toward the gate, and not a saving on the bill**: a workload that does not exist yet has no "before" on the Cost and Usage Report. It is in no total above.

Re-running this measurement

Everything here can be re-run by Copperwick's own team. The inputs are the CUR 2.0 export, the register of changes (the resource, tag, or ARN, and the date), and the cost expressions, the excluded line types, and the queries in [the method](#). To measure a change made after this report, wait for seven full days of finalized data, which with the 48-hour lag is nine days, run the query, multiply by 0.94, and add the row. The monthly cost review in the handover report keeps the register and re-runs the whole set each quarter, so the claim that the bill fell by this much can always be checked by someone who was not in the room.

© 2026 Vantalect, LLC. Sample deliverable; Copperwick Therapeutics, Inc. is fictional.

AWS, Amazon Web Services, AWS Graviton, and the AWS service names are trademarks of Amazon.com, Inc. or its affiliates. Oracle and Java are registered trademarks of Oracle Corporation and its affiliates. Microsoft, Windows, and SQL Server are trademarks of the Microsoft group of companies. Red Hat and Red Hat Enterprise Linux are trademarks of Red Hat, Inc. PostgreSQL is a registered trademark of the PostgreSQL Community Association of Canada. Kubernetes is a registered trademark of The Linux Foundation. Nextflow is a trademark of Seqera Labs. Vantalect is not affiliated with or endorsed by any of them. Copperwick Therapeutics, Inc. is fictional.